Tomographic Multi-Target Tracking and Its Relationship to PHD Filters

Roy L. Streit
Metron, Inc.
streit@metsci.com

Department of Electrical and Computer Engineering
University of Massachusetts – Dartmouth
r.streit@ieee.org

NATO Lecture Series
Multisensor Fusion: Advanced Methodologies and Applications

6 – 7 June 2011, Porto (PT)
9 – 10 June 2011, Paris (FR)
20 – 21 June 2011, Lancaster (UK)
23 – 24 June 2011, Ankara (TR)

"All journeys have secret destinations of which the traveler is unaware."
Martin Buber

- Understanding the PHD (Probability Hypothesis Density) target model
- "Secret" Destination
  - The connection between PHD filters and PET medical imaging

    PET = Positron Emission Tomography

Hidden Corollary

"The traveler is unaware of the direction of travel."

- Understanding a different kind of medical tomography called SPECT

  SPECT = Single Photon Emission Computed Tomography
- "Secret" Destination
  - The connection between SPECT and multi-sensor intensity filtering

Lecture 1

Topics – Lecture 1

- Background on Poisson point processes
- Single Sensor Multitarget Intensity / PHD filter
  - These are different, but essentially the same filters
- Method 1: Intuitive counting method
  - Reduces to the PHD filter by inserting prior knowledge
- Positron Emission Tomography (PET) imaging
  - PET is Poisson point process (PPP) intensity estimation
  - Shepp-Vardi algorithm (1984)
  - Identical (!) to Richardson-Lucy algorithm (1972 / 1974)
    - Different application to imaging & astronomy
- Method 2: First step of Shepp-Vardi algorithm
- Example: Multistatic sonar
  - Particle filter implementation
  - Simulated data
  - Movie

Topics – Lecture 2

- More examples
  - Augment kinematic data with feature data
  - Track extraction by the heritage method
  - More particle filter movies on simulated multistatic sonar data
- Method 3: First moment approximation
  - Mean field approximation and scaling
- Other topics
  - Regularization
  - Cramer-Rao Bounds (CRB)
- SPECT (Single Photon Emission Computed Tomography)
- Multiple Sensor Multi-Target Intensity Filtering
  - Variance reduction via multisensor filtering
- Concluding remarks

Poisson Point Processes (PPPs)

- PPPs are misunderstood models
  - Sometimes confused with the Poisson distribution
  - Often thought to be homogeneous
- Multidimensional and nonhomogeneous PPPs
  - Very flexible models
  - Easy to understand
  - Widely useful in applications
  - Main focus of interest here
- Tomography (PET and SPECT)
  - PPPs are excellent fit to the physics and the sensor data
- Multisensor and multitarget tracking/fusion
  - Model fit is less than perfect, but very useful approximation

The Poisson Distribution

- Random variable N is Poisson(λ) distributed if
  - It is defined on the non-negative integers: $\{0,\ 1,\ 2,\ldots\}$
  - The probability of $N = n$ is

$$p_N(n) \;=\; \Pr[N = n] \;=\; \exp(-\lambda)\,\frac{\lambda^n}{n!} \qquad \text{where } \lambda > 0$$

- This is a probability density (mass) function (pdf) since

$$\sum_{n=0}^{\infty} \Pr[N = n] \;=\; \sum_{n=0}^{\infty} p_N(n) = 1$$

- Important readily verified fact: $\mathrm{Var}[N] \;=\; \mathrm{Mean}[N] \;=\; \lambda$
- Rapid convergence of pdf to Gaussian shape:



Intensity Function

- Suppose there is a Poisson parameter λ for every point $x \in S \subset \mathbb{R}^m$
- Intensity function $\lambda(x) \geq 0$ is the "parameter" of a PPP
  - Knowing $\lambda(x)$ fully specifies the PPP
- Assumption:

$$\int_R \lambda(x)\, dx < \infty \quad \text{for every bounded subset } R \text{ of } S$$

  - R can be convex or concave, have holes, disconnected
  - It is not required that $\displaystyle\int_{R=\text{All } S} \lambda(x)\, dx < \infty$
- The PPP is non-homogeneous if $\lambda(x) \neq$ constant
  - The PPP is homogeneous if $\lambda(x) \equiv$ constant
- General form: $\lambda(x) = \lambda_{function}(x) + \sum_{k=1}^{\infty} w_k \delta(x - x_k), \quad w_k \geq 0$
  - $\lambda_{function}(x)$ can have step discontinuities
- Orderly PPPs: $\lambda(x)$ has no Dirac delta functions

Realizations in Two Easy Steps

- Realizations comprise all pairs $(n, \{x_1, \ldots, x_n\})$, $n \geq 0$, and $x_j \in R$ for all j
  - Set notation $\{x_1, \ldots, x_n\}$ emphasizes that ordering of points is irrelevant
  - Event space is collection of all pairs $\mathrm{E}(R) = \left\{ \left(n, \left\{x_1, \ldots, x_n\right\}\right) \right\}$
- Realizations are generated on bounded sets R by a two step procedure
- ***Step 1.*** Choose n with probability

$$p_N[n] = \exp\left(-\int_R \lambda(x)\, dx\right) \frac{\left(\int_R \lambda(x)\, dx\right)^n}{n!}$$

Poisson distribution

- ***Step 2.*** Choose n iid (independent and identically distributed) points with pdf

$$p_{X|N}(x) \equiv p_X(x) = \frac{\lambda(x)}{\int_R \lambda(s)\, ds} \quad \Rightarrow \quad \int_R p_X(x)\, dx = 1$$

- Realizations are lists
  - Points corresponding to the Dirac delta functions can be sampled repeatedly
- Orderly PPPs: Points are distinct with probability one

Seeing Detailed Structure from Samples is Nontrivial

Intensity function in 3-D

E[N]=100





Intensity function as a contour plot

Notch (*no points allowed*)

90 Samples

The notch is hard to see with only 90 samples





Samples superimposed on intensity function contours

[10]

Intensity and PHD Tracking Filters

Multi-Target Tracking

- Tracking in clutter is a joint detection and estimation problem
- Two inter-related tasks
  - Assignment of sensor measurements either to targets or to clutter
    - If assignments are known, the problem is “straightforward”
      - Track each target separately with clutter-free data
    - Assignment is the key computational roadblock
  - Initiation and termination of target tracks
- MHT (Multiple Hypothesis Tracking)
  - Widely accepted as an excellent model for many radar and sonar systems
  - Assignments are written as a list of “hypotheses”
    - Each hypothesis corresponds to a partition of the data
    - The components of a given partition correspond to individual targets or to clutter
  - Number of hypotheses grows exponentially with data and target counts
  - MHT implementations require pruning, and it therefore suboptimal
- Approximation is probably inevitable for any method
- Approximation assumed here: PPP model of the multitarget state

Modeling Gaussian Noise with PPP Intensity

- Bivariate i.i.d. Gaussian noise is modeled by the bivariate intensity function

$$\lambda(x,y) = c\ \mathrm{N}\left(\begin{bmatrix} x \\ y \end{bmatrix} \middle| \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix}\right)$$

where $c = \bar{N} = E[N] > 0$ is mean number of points in a PPP realization

$$\mathrm{N}\left(\vec{x} \mid \mu, \Sigma\right) = \frac{1}{\sqrt{\det(2\pi\Sigma)}} \exp\left(-\frac{1}{2}(\vec{x}-\mu)^T \Sigma^{-1} (\vec{x}-\mu)\right)$$

$=$ multivariate Gaussian pdf with mean μ and covariance Σ

- **Movie**: 25 scans with $c = 1$



$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Superposition: all 25 scans ($\Rightarrow c = 25$)



- Mean intensity $c = 1$ is not the same as one measurement per scan ($Pd = 1$)

$$\Pr[\text{ no } measurement] = \Pr[\text{ one } measurement] = e^{-1} = 37\%$$

PPP Model of Multitarget State

- *Interpretation #1*
  - The states of multiple targets are the points of a PPP realization
- *PPP realizations seriously mismodel multitarget state*
  - Example. Exactly one target is present with high Pd and in very low clutter, then

    Pr[target realization contains exactly one point] $= p_N(1) = e^{-1} \approx 0.37$

    Hence, 63% of the PPP realizations have no target, or more than one target
  - More generally: When exactly *n* targets are present, then

$$\begin{aligned}
&\Pr[\text{ target realization contains exactly } n \text{ points}] \\
&\equiv p_N(n) \\
&= e^{-n}\frac{n^n}{n!} \approx \frac{1}{\sqrt{2\pi n}} \to 0 \qquad (\text{Stirling's Approximation}) \\
&\approx \frac{1}{8} \quad \text{for } n = 10 \text{ targets}
\end{aligned}$$

- Hence, for large *n*, very few of the PPP realizations have the correct number of targets

PPPs for Multitarget Tracking

- *Interpretation #2*
  - Physical reality attaches to the concept of the *ensemble average* over all PPP realizations
  - The intensity function *parameterizes* the ensemble average
- In the language of thermodynamics, the points of a PPP realization are "microtargets"
- Why work so hard?
- PPP approximation finesses the data assignment problem
  - Key target measurement rule: "at most one measurement per target"
- *The assignment problem is replaced by an imaging problem*
  - Imaging problem is <u>much</u> easier to solve
- Huge benefit – no exponential growth in computational load

Augmented State Space

- Augment the state space S with a state ϕ that denotes "target absent"
  -- Appropriate space for joint detection and tracking (Kopec, IEEE ASAP, 1986)
- Integrals over the augmented space $S^+ = S \cup \phi$ are defined by

$$\int_{S^+} g(x)\, dx \equiv g(\phi) + \int_S g(x)\, dx$$

  where $g(\cdot)$ is a real valued function defined on S^+
- Measurement function is defined for all target states in S^+

  $p(z \mid x) =$ pdf of data z conditioned on target state x

  $p(z \mid \phi) =$ pdf of data z conditioned on target absent, ϕ
- Transition function (target motion model) defined on $S^+ \times S^+$

  $\Psi(y \mid x) = pdf$ [target moves from state x to state y at next time step]

  $\Psi(y \mid \phi) = pdf$ [target initiates in state y at next time step]

  $\Psi(\phi \mid x) =$ Pr[target in state x terminates at next time step]

  $\Psi(\phi \mid \phi) =$ Pr[absent target remains absent at next time step]

PPP Filter in Five Easy Pieces



Piece #1: Dynamics

- Multitarget PPP at time t_{k-1} has intensity function $f_{k-1|k-1}(x)$
- Target motion model is defined on the augmented space $S^+ = S \cup \phi$

$$\Psi_{k-1}(x \mid y) = \Pr\left[\text{target at } y \text{ at time } t_k \mid \text{target at } x \text{ at time } t_{k-1}\right]$$

- ***Piece 1***. Intensity of the "predicted" target PPP on S^+

$$\begin{aligned} f_{k|k-1}(x) &= \int_{S^+} \Psi_{k-1}(x \mid y)\, f_{k-1|k-1}(y)\, dy \\ &\equiv \underbrace{\Psi_{k-1}(x \mid \phi)\, f_{k-1|k-1}(\phi)}_{\text{Target birth process } b^{Birth}_{k|k-1}(x)} + \int_{S} \Psi_{k-1}(x \mid y)\, f_{k-1|k-1}(y)\, dy \end{aligned}$$

- The predicted target PPP intensity is defined on S^+
- Realizations can have repeated points with state ϕ

PPP Filter Block Diagram



The Piece With Finesse

- Input: Predicted target PPP at time t_k with intensity function $f_{k|k-1}(x)$
- $P_k^D(x)$ is probability of detecting a target at state x
  - $P_k^D(\phi) = \Pr\left[\text{absent target hypothesis generates a measurement}\right]$
- ***Piece #2***. Target detection (thinning)
  - *Detected* targets "pass" the detection test with probability $P_k^D(x)$

$$\boxed{f_{k|k-1}^D(x) = P_k^D(x)\, f_{k|k-1}(x)}$$

  - *Undetected* targets "fail" the detection test with probability $1 - P_k^D(x)$

$$\boxed{f_{k|k-1}^U(x) = \left(1 - P_k^D(x)\right) f_{k|k-1}(x)}$$

- Finesse: The "Detected" and "Undetected" target processes are *independent* PPPs
  - Fundamental to Poisson. Convincing Proof:
    - The number of heads H and the number of tails T in n coin flips are **independent** if the number of coin flips n is Poisson distributed

PPP Filter Block Diagram



Information-Updated Undetected Targets
-- The *Yogi Berra* Piece --

- ***Piece #3***. Information-updated undetected target PPP intensity

$$
\begin{aligned}
f^{U}_{k|k}(x) &= f^{U}_{k|k-1}(x) \\
&= f_{k|k-1}(x)\left(1 - P^{D}_{k}(x)\right)
\end{aligned}
$$

- Update is independent of the data
- "You can observe a lot just by watching."
  - Yogi Berra

PPP Filter Block Diagram



Intuitive Piece: Information-Update

- Given measurements are $\{z_1, \ldots, z_m\}$ at time t_k with pdf $p_k(z \mid x)$ for all $x \in S^+$
- Predicted detected target intensity is $f^D_{k|k-1}(x) = P^D_k(x)\, f_{k|k-1}(x)$
- Best current estimate

$$\Pr\left[z_j \text{ originates from ("}\leftarrow\text{") a target in } [x, x+dx)\right]$$

$$= \frac{E\left\{N\left[z_j \leftarrow [x, x+dx)\right]\right\}}{E\left\{N\left[z_j \leftarrow S^+\right]\right\}} = \frac{p_k(z_j \mid x)\, f^D_{k|k-1}(x)\ dx}{\int_{S^+} p_k(z_j \mid s)\, f^D_{k|k-1}(s)\, ds}$$

- "*At most one measurement per target*" rule $\Longrightarrow$

$$\sum_j \Pr\left[z_j \leftarrow [x, x+dx)\right]$$

$$= \text{ Estimated number of targets in } [x, x+dx) \text{ that generate measurements}$$

- $E\left[\text{Number of targets in } [x, x+dx)\right] = \int_x^{x+dx} f^D_{k|k}(x)dx \approx f^D_{k|k}(x)\ dx$
- Information-update of target rate:

$$\frac{\textit{Expected} \text{ number of targets}}{dx} = \frac{\textit{Estimated} \text{ number of targets}}{dx}$$

Information Update: Detected Targets

- Target rate per unit state space:

$$f^{D}_{k|k}(x) = f_{k|k-1}(x) \sum_{j=1}^{m} \frac{p_k(z_j \mid x)\, P^{D}_{k}(x)}{\int_{S^+} p_k(z_j \mid s)\, P^{D}_{k}(x)\, f_{k|k-1}(s)\, ds}$$

- This is the PPP *approximation*
- Bayesian derivation shows why it is an approximation
- Recursion is "closed"

  $\Leftrightarrow$ Initialized with a PPP, each iteration outputs a PPP

- PPP approximation is built into *every* iteration

PPP Filter Block Diagram



The Easy Piece: Superposition

- ***The 5th Piece***
  - Information updated detected and undetected target PPPs are independent
- Adding intensities gives final form of the multitarget PPP intensity filter:

$$f_{k|k}(x) = f^{D}_{k|k}(x) + f^{U}_{k|k}(x)$$

$$= f_{k|k-1}(x)\left(1 - P^{D}_{k}(x) + \sum_{j=1}^{m} \frac{p_k(z_j \mid x)\, P^{D}_{k}(x)}{\int_{S^+} p_k(z_j \mid s)\, P^{D}_{k}(s)\, f_{k|k-1}(s)\, ds}\right)$$

- Valid on the augmented space S^+
- Augmented space
  - Key to obtaining "*on line*" estimates of clutter and target birth intensities

iFilter

- Rewrite integral over the augmented state space S^+

$$\int_{S^+} p_k(z\,|\,s)\, P_D(s)\, f_{k|k-1}(s)\, ds$$

$$= \underbrace{\boxed{p_k(z\,|\,\phi)\, P_k^D(\phi)\, f_{k|k-1}(\phi)}}_{\lambda_{k|k-1}^{Clutter}(z)} + \underbrace{\boxed{\int_S p_k(z\,|\,s)\, P_k^D(s)\, f_{k|k-1}(s)\, ds}}_{\lambda_{k|k-1}^{Target}(z)}$$

Predicted *Clutter* Measurement Predicted *Target* Measurement

- Equivalent form:

$$f_{k|k}(x) = f_{k|k-1}(x)\left(1 - P_k^D(x) + \sum_{j=1}^{m} \frac{p_k(z_j\,|\,x)\, P_k^D(x)}{\lambda_{k|k-1}^{Clutter}(z_j) + \lambda_{k|k-1}^{Target}(z_j)}\right)$$

 PHD filter is obtained from iFilter

PHD Filter

- Limit filter to state space S using assumed *a priori* intensities
  - No transitions into ϕ (target death) $\Rightarrow$ thin $f_{k-1|k-1}(x)$ with death process
  - No transitions out of ϕ (target birth) $\Rightarrow$ substitute $b_{k|k-1}^{Birth}(x) \approx b_k^{Prior}(x)$
  - Clutter: $\lambda_{k|k-1}^{Clutter}(z) \approx \lambda_k^{Prior}(z)$
- These substitutions give the PHD filter

$$f_{k|k-1}(x) = b_k^{Prior}(x) + \int_S \Psi_k(x \mid y)\, f_{k-1|k-1}(y)\, dy$$

$$\lambda_{k|k-1}^{Target}(z) = \int_S p_k(z \mid s)\, P_k^D(s)\, f_{k|k-1}(s)\, ds$$

$$f_{k|k}(x) = f_{k|k-1}(x)\left(1 - P_D(x) + \sum_{j=1}^{m} \frac{p_k(z_j \mid x)\, P_k^D(x)}{\lambda_k^{Prior}(z_j) + \lambda_{k|k-1}^{Target}(z_j)}\right)$$

- Substitutions *a posteriori*
  - Standard but little used practice, e.g., Schmidt-Kalman filter (c. 1968)
  - May regularize filter if prior intensities are more reliable than estimates

Positron Emission Tomography

Connections to PHD and Intensity Filtering

Positron Emission Tomography (PET) Imaging

- PET imaging estimates intensity function of a PPP
  - Maximum *A Posteriori* (MAP) algorithm
- At least four independent discoveries of the algorithm
  - Computed tomography (Shepp-Vardi)
    - L. A. Shepp and Y. Vardi, "Maximum Likelihood Reconstruction for Emission Tomography," *IEEE Trans. Medical Imaging*, Vol. MI-1, 113-122, 1982
    - Lange and Carson, *J. Comp. Assisted Tomography*, 1984
  - Image deconvolution (Richardson-Lucy)
    - Richardson, *J. Optical Society of America*, 1972
    - Lucy, *The Astronomical Journal*, 1974
- All fielded PET medical imaging systems are variants of Shepp-Vardi
  - High resolution algorithms are practical
  - Mathematical models are well matched to physics and engineering practice

PET Imaging Physics

- Inject a short-lived radioisotope and wait 15 minutes or so
  - Radioisotope is attached to a sugar, e.g., fluorodesoxyglucose
  - Absorption depends on tissue type
    - Pre-cancerous cells typically show heightened metabolic activity
- Radioisotope undergoes beta decay
  - Each decay emits a single positron
- Positron encounters electron, and they annihilate
  - An annihilation event emits a pair of (gamma) photons
  - Photons move in opposite directions (if positrons are slow moving)
    - Conservation of momentum
- Raw “time of flight” data
  - Arrival times of photons at a bank of detectors
- Measurements are preprocessed raw data
  - Estimated locations of annihilation events

PET Imaging System



Detector "Stave" or Block



Source of images:
http://en.wikipedia.org/wiki/Positron_emission_tomography#Image_reconstruction (2003)

Typical PET Image of the Brain

"This is an image taken from a typical PET acquisition. It is a tomographic view of a brain examination in transaxial view. Red areas show more accumulated radioactivity and blue areas are partitions where low to no activity was accumulated. It should illustrate how a typical PET image looks like. It was taken with an ECAT Exact HR+ PET Scanner."

Source: http://en.wikipedia.org/wiki/Positron_emission_tomography#Image_reconstruction (2003)



PET Likelihood Function

- Let $R_1, \ldots, R_K$ be a close-packed disjoint pixel grid, and $\mathrm{R} = \bigcup R_k \subset \mathbb{R}^2$
  - Let λ_k denote the intensity of annihilation events in pixel R_k
- Let $\lambda(x) = \sum_{k=1}^{K} \lambda_k \, I_k(x)$ where $I_k(x) = 1$ if $x \in R_k$ and 0 otherwise
- Measured locations of annihilations form a realization of a PPP on $\mathbb{R}^2$
  - Data: $Z = \{z_1, \ldots, z_m\}$ where $z_j \in \mathbb{R}^2$
- Measurement likelihood function determined by the engineering system design

$$p(z \mid x) = pdf\left[\text{annihilation event detected at } z \;\middle|\; \text{event actually occurred at } x\right]$$

- Intensity of the measurement PPP is

$$\mu(z) = \int_{\mathrm{R}} p(z \mid x)\, \lambda(x)\, dx = \sum_{k=1}^{K} \lambda_k \, f_k(z) \quad \text{where} \quad f_k(z) = \int_{pixel\ R_k} p(z \mid x)\, dx$$

- The pdf of Z is

$$p(Z; \Lambda) = \exp\left(-\int_{\mathrm{R}} \mu(z)\, dz\right) \prod_{j=1}^{m} \mu(z_j)$$

- Substitute for $\mu(z)$

MAP Estimation

- MAP estimate

$$\hat{\Lambda}_{MAP} \equiv \arg\max_{\Lambda} \exp\left(-\sum_{k=1}^{K} \lambda_k \left|R_k\right|\right) \prod_{j=1}^{m} \left(\sum_{k=1}^{K} \lambda_k f_k(z_j)\right)$$

where $\Lambda \equiv (\lambda_1, \ldots, \lambda_K) \geq 0$

- **Method I.** Solve the necessary conditions $\nabla_{\Lambda} \mathrm{L}(Z;\Lambda) = 0$
  - L is the Lagrangian of $p(Z;\Lambda)$ and the constraints $\Lambda \geq 0$
  - Numerical details are unattractive
- **Method II.** Use Expectation-Maximization (EM)
  - Handles the product of sums very easily
  - Non-negativity constraints are automatically satisfied
  - Can be extended to incorporate other useful constraints

EM: Complete Data Likelihood

- EM method for PET (following Snyder and Miller, Chap. 3)
- Given: Measured locations of annihilation events $Z = \{z_1, \ldots, z_m\}$
- Unknown: True locations of annihilation events $X = \{x_1, \ldots, x_m\}$
- Missing data answer the question "Which pixel contains x_j?"

  $k_j \in \{1, \ldots, K\}$ is the index of the pixel containing x_j: $\boxed{x_j \in R_{k_j}}$
- Complete data (in the sense of EM)

  $\{(k_1, z_1), \ldots, (k_m, z_m)\}$
- Complete data pdf is

$$p(K, Z; \Lambda) = \exp\left(-\sum_{k=1}^{K} \lambda_k \left|R_k\right|\right) \prod_{j=1}^{m} \lambda_{k_j} \; f_{k_j}(z_j)$$

- Missing data change product of sums to product of summands

EM Method

- E-step:
  - Let $\Lambda^{(0)} = \left\{\lambda_1^{(0)}, \ldots, \lambda_K^{(0)}\right\} > 0$ be given
  - Let n = EM iteration index
  - Auxillary function

$$Q\left(\Lambda \mid \Lambda^{(n)}\right) = -\sum_{k=1}^{K} \lambda_k \underbrace{\left|R_k\right|}_{\text{pixel area}} + \sum_{k=1}^{K} \underbrace{\left(\sum_{j=1}^{m} \frac{\lambda_k^{(n)} f_k(z_j)}{\sum_{r=1}^{K} \lambda_r^{(n)} f_r(z_j)}\right)}_{\text{Known from previous EM iteration}} \log \lambda_k$$

- M-step:
  - Solve for $\Lambda^{(n+1)} = \arg\max_{\Lambda} Q\left(\Lambda \mid \Lambda^{(n)}\right)$

- The PET image is the parameter list $\Lambda = \lim_{n \to \infty} \Lambda^{(n+1)}$

Shepp-Vardi Algorithm

- Shepp-Vardi update for gridded images

$$\boxed{\lambda_k^{(n+1)} = \frac{\lambda_k^{(n)}}{|R_k|} \sum_{j=1}^{m} \left(\frac{f_k(z_j)}{\sum_{r=1}^{K} \lambda_r^{(n)} f_r(z_j)} \right)}$$

- Shepp-Vardi update for continuous images
  - Take limit as $|R_k| \to 0$ as number of pixels $K \to \infty$

$$\lambda_k^{(n)} \to \lambda^{(n)}(x)$$

$$\frac{f_k(z_j)}{|R_k|} = \frac{1}{|R_k|} \int_{R_k} p(z_j \mid x)\, dx \to p(z_j \mid x)$$

$$\sum_{r=1}^{K} \lambda_r^{(n)} f_r(z_j) \to \int_{\mathrm{R}} p(z_j \mid s)\, \lambda^{(n)}(s)\, ds$$

$$\boxed{\lambda^{(n+1)}(x) = \lambda^{(n)}(x) \sum_{j=1}^{m} \left(\frac{p(z_j \mid x)}{\int_{\mathrm{R}} p(z_j \mid s)\, \lambda^{(n)}(s)\, ds} \right)}$$

Same form as information update of detected target PPP.

Multistatic Sonar Example

- Simulated data
  - Generated by Pascal de Thieje, TNO, The Netherlands
  - 2008 for MSTWG (Multi-Static Tracking Working Group)
  - Heterogeneous sensors
  - Two source-receiver pairs
    - Each pair is a *de facto* sensor
  - One traditional monostatic sensor
  - Measurements at a sequence of specified times
    - Each comprises a set of data { (angle, bistatic range / range) }
    - Correspond to targets and clutter/interferers, but which is unknown
- Sensor data arrive at different times
  - Treated as a single sensor with time-varying characteristics
  - Exception:
    - Occasional time coincidence sensor reports treated as slightly different times

Multistatic Sonar Example

- Implementation by particle method
  - Based on the Nodestar single-target Bayesian filter
    - *Bayesian Multiple Target Tracking* by Stone, Barlow, and Corwin, 1999
  - Bryan Osborne, Kyrill Orloff, and Rebecca Wasyk (all Metron)
- Scenario
  - Sensor tracks (5 in all)
  - Target tracks are initiated manually
    - Automation coming soon (next lecture)
  - Point estimates via the "Mean Shift" algorithm
    - Widely used in image processing (Fukunaga, c. 1974)
    - EM method (Carreira-Perpinan, IEEE PAMI, 2007)
      - Maximum likelihood estimate
- Tracking movie generated by particle filter
  - 60 000 Monte Carlo particles
  - For details on particle filters, see *Beyond the Kalman Filter* by Ristic et al.

PhDFilter60KParticles.avi

More examples / movies in second lecture

Concluding Remarks – Lecture 1

- PHD and Intensity filters are *tomographic* imaging filters
  - Image space is target state space
- PET connection is insightful
  - Physical interpretation of the multiple target model
  - Convincing intuitive understanding of the information update
- Augmented space of the intensity filter (iFilter)
  - Provides on-line estimates of clutter PPPs and target birth
  - Prior knowledge reduces the intensity filter to the PHD filter
- Regularization
  - Reconstruction artifacts widely recognized in PET
    - Bayes: Use Markov Random Field as a prior for intensity function
  - Suggests the same need in PHD filter and iFilter
  - PET opens doors to 39 years of computational experience